



[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

Search: The ACM Digital Library The Guide

data type and external storage of large objects

SEARCH

THE ACM DIGITAL LIBRARY

[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used [data type](#) and [external storage of large objects](#)

Found 124,141 of 192,876

Sort results by

relevance

Save results to a Binder

Display results

expanded form

Search Tips

Open results in a new window

Try an [Advanced Search](#)

Try this search in [The ACM Guide](#)

Results 1 - 20 of 200

Result page: **1** [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale

1 [External memory algorithms and data structures: dealing with massive data](#)

Jeffrey Scott Vitter

June 2001 **ACM Computing Surveys (CSUR)**, Volume 33 Issue 2

Publisher: ACM Press

Full text available: [pdf\(828.46 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Data sets in large applications are often too massive to fit completely inside the computers internal memory. The resulting input/output communication (or I/O) between fast internal memory and slower external memory (such as disks) can be a major performance bottleneck. In this article we survey the state of the art in the design and analysis of external memory (or EM) algorithms and data structures, where the goal is to exploit locality in order to reduce the I/O costs. We consider a varie ...

Keywords: B-tree, I/O, batched, block, disk, dynamic, extendible hashing, external memory, hierarchical memory, multidimensional access methods, multilevel memory, online, out-of-core, secondary storage, sorting

2 [An open abstract-object storage system](#)

Stephen Blott, Lukas Relly, Hans-Jörg Schek

June 1996 **ACM SIGMOD Record , Proceedings of the 1996 ACM SIGMOD international conference on Management of data SIGMOD '96**, Volume 25 Issue 2

Publisher: ACM Press

Full text available: [pdf\(1.15 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Database systems must become more open to retain their relevance as a technology of choice and necessity. Openness implies not only databases exporting their data, but also exporting their services. This is as true in classical application areas as in non-classical (GIS, multimedia, design, etc).This paper addresses the problem of exporting storage-management services of indexing, replication and basic query processing. We describe an abstract-object storage model which provides the basic mechan ...

3 [Data base directions: the next steps](#)

John L. Berg

November 1976 **ACM SIGMOD Record , ACM SIGMIS Database**, Volume 8 , 8 Issue 4 , 2

Publisher: ACM Press

Full text available: [pdf\(9.95 MB\)](#) Additional Information: [full citation](#), [abstract](#)

What information about data base technology does a manager need to make prudent decisions about using this new technology? To provide this information the National Bureau of Standards and the Association for Computing Machinery established a workshop of approximately 80 experts in five major subject areas. The five subject areas were auditing, evolving technology, government regulations, standards, and user experience. Each area prepared a report contained in these proceedings. The proceedings p ...

Keywords: DBMS, auditing, cost/benefit analysis, data base, data base management, government regulation, management objectives, privacy, security, standards, technology assessment, user experience

4 Enhanced abstract data types in object-relational databases

Praveen Seshadri

August 1998 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 7 Issue 3

Publisher: Springer-Verlag New York, Inc.

Full text available: [pdf\(119.21 KB\)](#) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

The explosion in complex multimedia content makes it crucial for database systems to support such data efficiently. This paper argues that the "blackbox" ADTs used in current object-relational systems inhibit their performance, thereby limiting their use in emerging applications. Instead, the next generation of object-relational database systems should be based on enhanced abstract data type (E-ADT) technology. An (E-ADT) can expose the semantics of its methods to the database ...

Keywords: Database types, Extensibility, Object-relational database, Query optimization

5 A Value Transmission Method for Abstract Data Types

Maurice P. Herlihy, Barbara Liskov

October 1982 **ACM Transactions on Programming Languages and Systems (TOPLAS)**, Volume 4 Issue 4

Publisher: ACM Press

Full text available: [pdf\(1.63 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

6 Persistent memory: a storage architecture for object-oriented database systems

Satish M. Thatte

September 1986 **Proceedings on the 1986 international workshop on Object-oriented database systems**

Publisher: IEEE Computer Society Press

Full text available: [pdf\(1.13 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Object-oriented databases are needed to support database objects with a wide variety of types and structures. A persistent memory system provides a storage architecture for long-term, reliable retention of objects with rich types and structures in the virtual memory itself. It is based on a uniform memory abstraction, which eliminates the distinction between transient objects (data structures) and persistent objects (files and databases), and therefore, allows the same set of powerful and f ...

7 Query evaluation techniques for large databases

Goetz Graefe

June 1993 **ACM Computing Surveys (CSUR)**, Volume 25 Issue 2



Publisher: ACM Press

Full text available: [pdf\(9.37 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Database management systems will continue to manage large data volumes. Thus, efficient algorithms for accessing and manipulating large sets and sequences will be required to provide acceptable performance. The advent of object-oriented and extensible database systems will not solve this problem. On the contrary, modern data models exacerbate the problem: In order to manipulate large sets of complex objects as efficiently as today's database systems manipulate simple records, query-processi ...

Keywords: complex query evaluation plans, dynamic query evaluation plans, extensible database systems, iterators, object-oriented database systems, operator model of parallelization, parallel algorithms, relational database systems, set-matching algorithms, sort-hash duality

8 Data-centric storage in sensornets with GHT, a geographic hash table

Sylvia Ratnasamy, Brad Karp, Scott Shenker, Deborah Estrin, Ramesh Govindan, Li Yin, Fang Yu

August 2003 **Mobile Networks and Applications**, Volume 8 Issue 4

Publisher: Kluwer Academic Publishers

Full text available: [pdf\(255.10 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Making effective use of the vast amounts of data gathered by large-scale sensor networks (sensornets) will require scalable, self-organizing, and energy-efficient data dissemination algorithms. For sensornets, where the content of the data is more important than the identity of the node that gathers them, researchers have found it useful to move away from the Internet's point-to-point communication abstraction and instead adopt abstractions that are more data-centric. This approach entails na ...

Keywords: algorithms, distributed systems, performance, sensor networks

9 Design of the Mneme persistent object store

J. Eliot B. Moss

April 1990 **ACM Transactions on Information Systems (TOIS)**, Volume 8 Issue 2

Publisher: ACM Press

Full text available: [pdf\(3.22 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

The Mneme project is an investigation of techniques for integrating programming language and database features to provide better support for cooperative, information-intensive tasks such as computer-aided software engineering. The project strategy is to implement efficient, distributed, persistent programming languages. We report here on the Mneme persistent object store, a fundamental component of the project, discussing its design and initial prototype. Mneme stores objects

10 A distributed repository for immutable persistent objects

Douglas Wiebe

June 1986 **ACM SIGPLAN Notices , Conference proceedings on Object-oriented programming systems, languages and applications OOPSLA '86**, Volume 21 Issue 11

Publisher: ACM Press

Full text available: [pdf\(1.00 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Jasmine is an object-oriented system for programming-in-the-large. Jasmine describes software using system model objects. These objects are persistent (they have lifetimes of days or decades) and immutable (since system models act as historical records). This paper describes JStore, a distributed, replicated repository for system model objects. JStore provides robust, transactional, write-once storage. Designs are presented for the serialization, ...

11 A storage system for complex objects

U. Deppisch, H.-B. Paul, H.-J. Schek

September 1986 **Proceedings on the 1986 international workshop on Object-oriented database systems**

Publisher: IEEE Computer Society Press

Full text available:  pdf(1.17 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Complex objects are required in many new applications of databases. A common characteristic feature is that objects use other (sub-) objects for their description. Consequently retrieval or extraction of complex objects may include some or all of their subobjects which - in turn - may have subobjects to be extracted too. Accordingly a storage system is described which was designed and implemented with the objective to provide this set orientation: Relations with relation-valued attr ...

12 Indexing for data models with constraints and classes (extended abstract)

 Paris C. Kanellakis, Sridhar Ramaswamy, Darren E. Vengroff, Jeffrey S. Vitter

August 1993 **Proceedings of the twelfth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems**

Publisher: ACM Press

Full text available:  pdf(1.08 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We examine I/O-efficient data structures that provide indexing support for new data models. The database languages of these models include concepts from constraint programming (e.g., relational tuples are generalized to conjunctions of constraints) and from object-oriented programming (e.g., objects are organized in class hierarchies). Let n be the size of the database, c the number of classes, B the secondary storage page size, and

13 Applications and OS: GHT: a geographic hash table for data-centric storage

 Sylvia Ratnasamy, Brad Karp, Li Yin, Fang Yu, Deborah Estrin, Ramesh Govindan, Scott Shenker

September 2002 **Proceedings of the 1st ACM international workshop on Wireless sensor networks and applications**

Publisher: ACM Press

Full text available:  pdf(217.28 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Making effective use of the vast amounts of data gathered by large-scale sensor networks will require scalable, self-organizing, and energy-efficient data dissemination algorithms. Previous work has identified data-centric routing as one such method. In an associated position paper [23], we argue that a companion method, data-centric storage (DCS), is also a useful approach. Under DCS, sensed data are stored at a node determined by the name associated with the sensed data. In this paper, we des ...

14 A survey of structured and object-oriented software specification methods and techniques

 Roel Wieringa

December 1998 **ACM Computing Surveys (CSUR)**, Volume 30 Issue 4

Publisher: ACM Press

Full text available:  pdf(605.26 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

This article surveys techniques used in structured and object-oriented software specification methods. The techniques are classified as techniques for the specification of external interaction and internal decomposition. The external specification techniques are further subdivided into techniques for the specification of functions, behavior, and communication. After surveying the techniques, we summarize the way they are used in structured and object-oriented methods and indicate ways in w ...

Keywords: languages

15 Special issue on persistent object systems: Orthogonally persistent object systems 

Malcolm Atkinson, Ronald Morrison

July 1995 **The VLDB Journal — The International Journal on Very Large Data Bases**,
Volume 4 Issue 3

Publisher: Springer-Verlag New York, Inc.

Full text available:  pdf(5.02 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

Persistent Application Systems (PASs) are of increasing social and economic importance. They have the potential to be long-lived, concurrently accessed, and consist of large bodies of data and programs. Typical examples of PASs are CAD/CAM systems, office automation, CASE tools, software engineering environments, and patient-care support systems in hospitals. Orthogonally persistent object systems are intended to provide improved support for the design, construction, maintenance, and operation o ...

Keywords: database programming languages, orthogonal persistence, persistent application systems, persistent programming languages

16 Extensions to Starburst: objects, types, functions, and rules 

 Guy M. Lohman, Bruce Lindsay, Hamid Pirahesh, K. Bernhard Schiefer
October 1991 **Communications of the ACM**, Volume 34 Issue 10

Publisher: ACM Press

Full text available:  pdf(5.21 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

Keywords: Extended relational database management systems, Starburst, extensible database management systems

17 A shared, segmented memory system for an object-oriented database 

 Mark F. Hornick, Stanley B. Zdonik
January 1987 **ACM Transactions on Information Systems (TOIS)**, Volume 5 Issue 1

Publisher: ACM Press

Full text available:  pdf(2.05 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

This paper describes the basic data model of an object-oriented database and the basic architecture of the system implementing it. In particular, a secondary storage segmentation scheme and a transaction-processing scheme are discussed. The segmentation scheme allows for arbitrary clustering of objects, including duplicates. The transaction scheme allows for many different sharing protocols ranging from those that enforce serializability to those that are nonserializable and require communi ...

18 Types and persistence in database programming languages Malcolm P. Atkinson, O. Peter BunemanJune 1987 **ACM Computing Surveys (CSUR)**, Volume 19 Issue 2**Publisher:** ACM PressFull text available:  pdf(7.91 MB)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Traditionally, the interface between a programming language and a database has either been through a set of relatively low-level subroutine calls, or it has required some form of embedding of one language in another. Recently, the necessity of integrating database and programming language techniques has received some long-overdue recognition. In response, a number of attempts have been made to construct programming languages with completely integrated database management systems. These lang ...

19 Strategic directions in storage I/O issues in large-scale computing Garth A. Gibson, Jeffrey Scott Vitter, John WilkesDecember 1996 **ACM Computing Surveys (CSUR)**, Volume 28 Issue 4**Publisher:** ACM PressFull text available:  pdf(465.35 KB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)**20 Fortran 8X draft** Loren P. MeissnerDecember 1989 **ACM SIGPLAN Fortran Forum**, Volume 8 Issue 4**Publisher:** ACM PressFull text available:  pdf(21.36 MB) Additional Information: [full citation](#), [abstract](#), [index terms](#)

Standard Programming Language Fortran. This standard specifies the form and establishes the interpretation of programs expressed in the Fortran language. It consists of the specification of the language Fortran. No subsets are specified in this standard. The previous standard, commonly known as "FORTRAN 77", is entirely contained within this standard, known as "Fortran 8x". Therefore, any standard-conforming FORTRAN 77 program is standard conforming under this standard. New features can b ...

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)